

METHOD AND SYSTEM FOR PROVIDING
A PROTECTION PATH FOR CONNECTIONLESS SIGNALS
IN A TELECOMMUNICATIONS NETWORK

RELATED APPLICATIONS

5 This application claims the benefit of U.S.
Provisional Application Serial No. 60/202,190, entitled
INTERNET PROTOCOL TRANSPORT, filed May 5, 2000 which is
hereby incorporated by reference.

TECHNICAL FIELD OF THE INVENTION

10 This invention relates generally to the field of
telecommunications and more particularly to a method and
system for providing a protection path for connectionless
signals in a telecommunications network.

BACKGROUND OF THE INVENTION

Telecommunications systems generally operate in either a connection-oriented mode or a connectionless mode. In a connectionless mode of operation, signals are communicated with less regard for the particular path traversed between source and destination and network elements than in a connection-oriented mode. Connectionless signaling typically focuses on the destination address, or other identification, rather than any particular path between source and destination network elements. Internet Protocol (IP), IPx, and SNA packet switching are examples of connectionless signal transport.

When a failure occurs along the working path being traversed by signals in connectionless communication, the signals must be re-routed to the destination network element along another available path. In conventional telecommunications systems, this re-routing is done by each individual network element with no pre-defined protection paths existing for the connectionless signals. Thus, these systems are inefficient. In addition, bandwidth is wasted by these systems due to inefficient bandwidth reservation schemes.

SUMMARY OF THE INVENTION

In accordance with the present invention, a method and system for providing a protection path for connectionless signals in a telecommunications network are provided that substantially eliminate or reduce disadvantages and problems associated with previously developed systems and methods. In particular, protection paths are pre-defined for connectionless signals, thereby increasing network efficiency.

In one embodiment of the present invention, a method is provided for providing protection for connectionless signals in a telecommunications network comprising a plurality of nodes. A first protection path is generated from each of the nodes to a destination node. A second protection path is generated from each of the nodes to the destination node. The second protection path is distinct from the first protection path. Protection traffic is routed along one of the protection paths to the destination node.

In another embodiment of the present invention, a node is provided in a telecommunications network. The node includes at least two ports and a protection egress port identifier. Each of the ports is operable to receive and transmit traffic for the node. The protection egress port identifier is operable to identify one of the ports as a protection egress port for a specified ingress port and a specified destination node. The protection egress port is operable to transmit protection traffic received at the specified ingress port for the specified destination node.

Technical advantages of the present invention include providing a method for providing a protection path for

connectionless signals in a telecommunications network. In particular, two protection paths are provided from each node to each destination node. Accordingly, each node may communicate with each other node along two distinct
5 protection paths. As a result, the network is protected against a single failure. In addition, network efficiency is improved through the use of the pre-defined protection paths.

Other technical advantages will be readily apparent to
10 one skilled in the art from the following figures, description, and claims.

009090" 02062550

BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the present invention and its advantages, reference is now made to the following description taken in conjunction with the accompanying drawings, wherein like numerals represent like parts, in which:

FIGURE 1 is a block diagram illustrating a communication system operable to facilitate communication of connectionless signals in accordance with one embodiment of the present invention;

FIGURE 2 is a block diagram illustrating a system for providing protection for connectionless signals communicated between the nodes of FIGURE 1 in accordance with one embodiment of the present invention;

FIGURE 3 is a block diagram illustrating a system for generating the protection paths of FIGURE 2 in accordance with one embodiment of the present invention;

FIGURE 4 is a block diagram illustrating one of the nodes of FIGURES 1, 2 and 3 operable to provide protection for connectionless signals in accordance with one embodiment of the present invention;

FIGURE 5 is a flow diagram illustrating a method for assigning a secondary protection egress port for the node of FIGURE 4 for each of a plurality of destination nodes;

FIGURE 6 is a flow diagram illustrating a method for assigning a protection egress port for each port for the node of FIGURE 4 for each of a plurality of destination nodes; and

ATTORNEY DOCKET NUMBER
064731.0142

PATENT APPLICATION

6

FIGURE 7 is a flow diagram illustrating a method for reserving bandwidth for connectionless signals communicated between the nodes of FIGURE 1 in accordance with one embodiment of the present invention.

09589038 060600

DETAILED DESCRIPTION OF THE INVENTION

FIGURE 1 is a block diagram illustrating a communication system 10 operable to facilitate communication of connectionless signals in accordance with one embodiment of the present invention. The telecommunications network is a network that transmits voice, audio, video or other suitable types of information, and/or a combination of different types of information between source and destination points. As used herein, the term "connectionless signal" refers to a signal that is not necessarily associated with any particular path from a source network element to a destination network element. In connectionless signaling, routing determinations, such as a determination of the next network element in the path, are generally made at each node in the path. Thus, each node independently identifies the next node in the forwarding chain. Connectionless signals include, for example, Internet Protocol (IP), IPx, SNA and other packet-switched transport signals.

The system 10 is operable to provide either connectionless communication or a hybrid of connectionless and connection-oriented communication, as described in the co-owned U.S. Patent Application entitled, "System and Method for Connectionless/Connection Oriented Signal Transport," filed on June 6, 2000. The system 10 comprises a core cloud 12 that comprises one or more core network elements 14, or nodes 14. The nodes 14 may communicate with each other via communication links 16 and with one or more peripheral network elements 20 via communication links 30. The communication links 16 and 30 may comprise any wireless, wireline, fiber or other communication medium or

combinations of media. A signal communicated via communication links 16 and/or 30 may comprise an electrical signal, an optical signal, or any other suitable type of signal or combination of signals.

5 The peripheral network elements 20 facilitate communication between the core cloud 12 and other network elements coupled to other networks, such as networks 36. According to the illustrated embodiment, the peripheral network elements 20 comprise routers 20. Each router 20
10 couples the core cloud 12 to a network 36 via a communication link 50. As used herein, "each" means every one of at least a subset of the identified items.

The routers 20 facilitate routing functions for signals originated or forwarded by interface equipment 40
15 and communicated over the networks 36. The interface units 40 comprise personal computers, servers, switches, routers or any other suitable network equipment operable to originate or forward communication signals.

The networks 36 may comprise any suitable wireline or
20 wireless systems that support communication between network elements using ground-based and/or space-based components. For example, the networks 36 may comprise public switched telephone networks, integrated services digital networks, local area networks, wide area networks, or any other
25 suitable communication system or combination of communication systems at one or more locations. Each of the networks 36 may comprise a single network or multiple networks.

In operation, the core cloud 12 receives
30 connectionless signals from the routers 20 and routes those signals through the core cloud 12 to another appropriate

5

10

15

25

traffic along the blue protection path 100 and a red port for transmitting traffic along the red protection path 102.

As illustrated in FIGURE 2, a destination node 108 may be reached by any other node 14 by following either the blue protection path 100 or the red protection path 102 from the other node 14.

In order to provide protection with these protection paths 100 and 102, a network must be protectable. The following notation will be used in determining whether or not a network is protectable. A graph G with n vertices, or nodes, and m edges, or links, has a vertex set $V(G)$ and an edge set $E(G)$. The count of the set $V(G)$ is expressed as $C(V(G))$, which is equivalent to the number of nodes 14 inside the network. A network is protectable against a single fault if for all x, y in the set G there are two distinct paths inside G . This is the case for any network for which an ear decomposition, as described below, is possible.

A path addition to graph G is the addition to G of a path of length $l \geq 1$ between two vertices of G , introducing $l-1$ new vertices. This added path is referred to as an ear. An ear decomposition is a partition of the graph G into sets P_0, \dots, P_k such that $C=P_0$ is a cycle and P_i for $i-1$ is an ear, or path addition, to the graph G formed by $P_0 \cup \dots \cup P_{i-1}$.

Thus, for each node 14 in a protectable network, there are two distinct paths 100 and 102 to reach any other node 14 in the network. The paths 100 and 102 are distinct in that they share no common nodes 14 or links. Using the ear decomposition previously described, the protectable network may be decomposed into a ring 120 and a set of ears 130.

In accordance with one embodiment, each node 14 in the network performs the ear decomposition in the same manner to arrive at the same result. It will be understood, however, that any suitable decomposition of the network may
5 be performed by the nodes to arrive at a same result without departing from the scope of the present invention.

The ring 120 forms the initial part of the protectable network and is denoted by P_0 . The first ear 130a comprises a set of nodes 14 coupled in a linear fashion with two
10 links coupled to two different nodes 14 in the ring 120. The first ear 130a is denoted as P_1 . The second ear 130b is a set of nodes 14 coupled to two different nodes 14 in either P_0 or P_1 . Although the illustrated embodiment comprises three ears 130a, b and c, it will be understood
15 that an ear decomposition of a protectable network may result in any suitable number of ears 130.

FIGURE 3 is a block diagram illustrating a system 140 for generating the protection paths 100 and 102 for connectionless signals communicated between the nodes 14 in
20 accordance with one embodiment of the present invention. After a protectable network has been decomposed through an ear decomposition, the protection paths 100 and 102 are generated as described in connection with FIGURE 3.

The ring 120 is split open at the destination node 108 and charted horizontally, beginning and ending with the
25 same destination node 108. Beginning with the leftmost destination node 108, the first ear 130d is identified based on a minimum hop count to the destination node 108. This ear 130d is denoted as P_1 and is charted horizontally
30 above the ring 120.

Any other ears 130, such as ears 130e, f and g, are identified in a similar manner and charted horizontally above the previous ear 130. After charting the ring 120 and the ears 130, the blue protection path 100 is generated from right to left on the chart and the red protection path 102 is generated from left to right on the chart. Thus, from any node 14, the protection paths 100 and 102 to the destination node 108 do not intersect, resulting in a network which is protected from any single failure.

FIGURE 4 is a block diagram illustrating one of the nodes 14 operable to provide protection for connectionless signals in accordance with one embodiment of the present invention. The node 14 comprises a plurality of ports 200, each of which is operable to receive traffic from and transmit traffic to other nodes 14 in the network. Although the illustrated embodiment comprises seven ports 200, it will be understood that any suitable number of ports 200 greater than one may be implemented in the node 14 without departing from the scope of the present invention.

The node 14 also comprises a traffic classifier 206 for classifying traffic received through ports 200 as either working traffic or protection traffic, a working traffic egress port identifier 210 for identifying an egress port 200 for working traffic, a protection egress port identifier 214 for identifying an egress port 200 for protection traffic, a secondary protection egress port identifier 218 for identifying an egress port 200 for protection traffic, an egress port evaluator 224 for evaluating the status of an egress port 200, and an egress

port selector 228 for selecting an appropriate egress port 200 and, in certain situations, discarding traffic.

5 The traffic classifier 206 is operable to classify traffic received through the ports 200 as either working traffic or protection traffic. According to one embodiment, the traffic comprises a traffic identifier identifying itself as either working traffic or protection traffic. For example, the traffic identifier may comprise a specified bit in the traffic, with one value for the bit
10 indicating working traffic and another value for the bit indicating protection traffic. The traffic classifier 206 is then operable to classify the received traffic based on the traffic identifier for the traffic. It will be understood that the traffic classifier 206 may classify the
15 traffic as working or protection traffic in any other suitable manner without departing from the scope of the present invention.

20 The working traffic egress port identifier 210 identifies an egress port 200 for working traffic received at any one of the ports 200 for the node 14 based on the corresponding destination node for the traffic. Thus, the working traffic egress port identifier 210 may comprise a database, table or other suitable data store for identifying a particular port 200 as an egress port for
25 working traffic received at the node 14 based on the destination node for the working traffic. For example, for working traffic received at the node 14 having a particular destination node, the working traffic egress port identifier 210 may identify port F 200 as the egress port
30 200 for that traffic. For another destination node, the

working traffic egress port identifier 210 may identify port C 200 as the egress port 200.

5 The protection egress port identifier 214 is operable to identify an egress port 200 for protection traffic received at any other port 200 in the node 14 bound for a particular destination node. Thus, the protection egress port identifier 214 may comprise a database, table or other suitable data store for identifying a particular port 200 as an egress port for protection traffic received at the
10 node 14 based on a particular ingress port 200 in conjunction with the destination node for the protection traffic. For example, for protection traffic received at port B 200 having a particular destination node, the protection traffic egress port identifier 214 may identify
15 port D 200 as the egress port 200 for that traffic. For protection traffic received at port B 200 having a different destination node, the protection traffic egress port identifier 214 may identify port E 200 as the egress port 200 for that traffic.

20 The secondary protection egress port identifier 218 is operable to identify an egress port 200 for transmitting traffic received at another port 200 in the node 14 for a particular destination node. Thus, the secondary protection egress port identifier 218 may comprise a
25 database, table or other suitable data store for identifying a particular port 200 as an egress port for transmitting protection traffic based on working traffic received at the node 14 having a particular destination node. For example, for working traffic received at the
30 node 14 having a particular destination node and for which the working traffic egress port 200 identified by the

working traffic egress port identifier 210 is unavailable, the secondary protection egress port identifier 218 may identify port G 200 as the egress port 200 for that traffic. Thus, for a particular destination node, the
5 secondary protection egress port identifier 218 identifies one egress port 200.

The egress port evaluator 224 is operable to evaluate the status of an egress port 200 in order to determine whether the port 200 is available or unavailable. The
10 status for a port 200 is available when the port 200 is functioning properly, while the status is unavailable when the port 200 is not functioning properly. According to one embodiment, each port 200 for the node 14 provides a status to the egress port evaluator 224. Alternatively, the
15 egress port evaluator 224 may test the ports 200 to determine status or may request a status from the ports 200. It will be understood that the egress port evaluator 224 may obtain the status of the egress ports 200 in any suitable manner without departing from the scope of the
20 present invention.

The egress port selector 228 is operable to select an appropriate egress port 200 and to discard traffic for which no appropriate egress port 200 is available. Thus, the egress port selector 228 may retrieve information from
25 the traffic classifier 206, the egress port evaluator 224, and one or more of the working traffic egress port identifier 210, protection egress port identifier 214, and secondary protection egress port identifier 218 in order to select the appropriate egress port 200 for a given
30 situation.

In operation, the traffic classifier 206, the working traffic egress port identifier 210, the protection egress port identifier 214, the secondary protection egress port identifier 218, the egress port evaluator 224 and the egress port selector 228 operate together to route traffic from a particular ingress port 200 to the appropriate egress port 200 for the node 14. The working traffic egress port 200 identified by the working traffic egress port identifier 210 is selected by the egress port selector 228 for received working traffic that is being forwarded along the working path 104. The protection egress port 200 identified by the protection egress port identifier 214 is selected by the egress port selector 228 for received protection traffic that is being forwarded along a protection path 100 or 102. The secondary protection egress port 200 identified by the secondary protection egress port identifier 218 is selected by the egress port selector 228 for traffic received on the working path 104, but that is being transmitted onto a protection path 100 or 102.

The egress port selector 228 is operable to select an egress port 200 based on traffic classification and on port status. If the traffic classifier 206 has classified the received traffic as protection traffic, the egress port selector 228 provides the protection egress port 200 identified by the protection egress port identifier 214 to the egress port evaluator 224 for status evaluation. If the protection egress port status is available, the egress port selector 228 selects the protection egress port 200 as the egress port 200 for the corresponding traffic. However, if the protection egress port status is

unavailable, the egress port selector 228 discards the corresponding traffic.

09589038 060600

If the traffic classifier 206 has classified the received traffic as working traffic, the egress port selector 228 provides the working traffic egress port 200 identified by the working traffic egress port identifier 210 to the egress port evaluator 224 for status evaluation. If the working traffic egress port status is available, the egress port selector 228 selects the working traffic egress port 200 as the egress port 200 for the corresponding traffic. However, if the working traffic egress port status is unavailable, the secondary protection egress port 200 identified by the secondary protection egress port identifier 218 is provided by the egress port selector 228 to the egress port evaluator 224 for status evaluation.

If the secondary protection egress port status is available, the egress port selector 228 selects the secondary protection egress port 200 as the egress port 200 for the corresponding traffic. However, if the secondary protection egress port status is unavailable, the egress port selector 228 discards the corresponding traffic.

Thus, for received working traffic with an available corresponding working traffic egress port 200, the traffic is transmitted by the working traffic egress port 200 along the working path 104. Similarly, for protection traffic with an available protection egress port 200, the traffic is transmitted by the protection egress port 200 along the protection path 100 or 102 on which the traffic was received. For received working traffic with an unavailable working traffic egress port 200, the traffic is transmitted onto a protection path 100 or 102 by the secondary

protection egress port 200, if available. Thus, this traffic is re-routed from the working path 104 to the protection path 100 or 102 having the shortest distance to the destination node. In addition, the traffic identifier
5 for the traffic is changed to indicate that the traffic is now protection traffic as opposed to working traffic.

FIGURE 5 is a flow diagram illustrating a method for assigning a secondary protection egress port for a particular node 14, referred to as node β , for each of a
10 plurality of destination nodes 64. The method begins at step 500 where a selection is made of a destination node 64 whose secondary protection egress port for node β has not been assigned. At step 501, a cost is determined for the blue protection path 100 based on the distance from the
15 blue port for node β to the destination node. At step 502, a cost is determined for the red protection path 102 based on the distance from the red port for node β to the destination node. At step 504 the cost from the working traffic egress port 200 to the destination node is set to
20 infinity. Thus, for the situation in which either the blue port or the red port comprises the working traffic egress port 200, the cost for the corresponding protection path 100 or 102 is infinity.

At step 506, the costs of the protection paths 100 and
25 102 are compared. At decisional step 508, a determination is made regarding whether the cost of the blue protection path 100 is less than the cost of the red protection path 102. If the cost of the blue protection path 100 is less than the cost of the red protection path 102, the method
30 follows the Yes branch from decisional step 508 to step 510. At step 510, the secondary protection egress port 200

for node β for the selected destination node is set to the blue port.

Returning to decisional step 508, if the cost of the blue protection path 100 is not less than the cost of the red protection path 102, the method follows the No branch from decisional step 508 to step 512. At step 512, the secondary protection egress port 200 for node β for the selected destination node is set to the red port.

From steps 510 and 512, the method continues to decisional step 514. At decisional step 514, a determination is made regarding whether any destination nodes 64 remain whose secondary protection egress port for node β has not been assigned. If destination nodes 64 remain whose secondary protection egress port for node β has not been assigned, the method follows the Yes branch from decisional step 514 and returns to step 500 for the selection of another destination node 64. However, if no destination node 64 remains whose secondary protection egress port for node β has not been assigned, the method follows No branch from decisional step 514 and comes to an end.

FIGURE 6 is a flow diagram illustrating a method for assigning a protection egress port for each port 200 for a particular node 14, referred to as node β , for each of a plurality of destination nodes 64. The method begins at step 600 where a selection is made of a destination node 64 whose protection egress ports for node β have not been assigned. This destination node 64 is denoted as d . At step 602, the ring 120 of the ear decomposition, P_0 , is identified by determining which ring (i) has the largest

nodal count, $C(V(P_0))$, among all the possible decompositions and (ii) has the largest node ID for the nodes 14 in $\overline{\{d\}} \cap V(P_0)$. Each node 14 has a corresponding node ID. The largest node ID is used to select a single ring 120 from
5 multiple possible rings which each have the same nodal count. The ring 120 starts and ends with the destination node d 64.

At step 604, the ring 120 is charted horizontally on a graph. Because the destination node d 64 is included
10 twice, at the beginning and end of the ring 120, the destination node d 64 has two coordinates on the graph. The left coordinate, which is expressed as $X(d')$, is zero, while the right coordinate, which is expressed as $X(d'')$, is α (a value greater than zero).

15 At step 606, each node 14 in the ring 120 is given a coordinate. Starting from the d' side, the j^{th} node in P_0 , denoted as $n_{P_0,j}$ will have a coordinate given by:

$$X(n_{P_0,j}) = \frac{j \times \alpha}{C(P_0)} \text{ with } j=0, \dots, C(P_0)-1 \quad (\text{eqn.1})$$

20

At decisional step 608, a determination is made regarding whether node β is in the ring 120. If node β is in the ring 120, the method follows the Yes branch from decisional step 608 to step 610. At step 610, a blue port
25 for node β is defined as the port 200 linked to node $n_{P_0,b} \in V(P_0)$ with coordinate $X(n_{P_0,b}) < X(\beta)$. At step 612, a red port for node β is defined as the port 200 linked to node

$n_{p_{0,r}} \in V(P_0)$ with coordinate $X(n_{p_{0,r}}) > X(\beta)$. At step 614, the protection egress port for the red port is set to the blue port. At step 616, the protection egress port for the blue port is set to the red port.

5 At decisional step 618, a determination is made regarding whether each port 200 of node β has been assigned a corresponding protection egress port for the destination node d 64. If each port 200 of node β has been assigned a corresponding protection egress port for the destination
10 node d 64, the method follows the Yes branch from decisional step 618 to decisional step 620.

At decisional step 620, a determination is made regarding whether any destination nodes d 64 remain whose protection egress ports for node β have not been assigned.
15 If destination nodes d 64 remain whose protection egress ports for node β have not been assigned, the method follows the Yes branch from decisional step 620 and returns to step 600 for the selection of another destination node d 64. However, if no destination nodes D 64 remain whose
20 protection egress ports for node β have not been assigned, the method follows the No branch from decisional step 620 and comes to an end.

Returning to decisional step 618, if each port of node β has not been assigned a corresponding protection egress
25 port for the destination node d 64, the method follows the No branch from decisional step 618 to step 619. At step 619, protection egress ports are assigned for the remaining ports 200 of node β .

Returning to decisional step 608, if node β is not in the ring 120, the method follows the No branch from decisional step 608 to step 622. At step 622, a counter, i , is set to zero. At step 624, the counter, i , is incremented by one.

At step 626, the i^{th} ear 130 of the ear decomposition, P_i , is identified by determining which ear (i) has the largest nodal count, $C(V(P_i))$, among all the possible decompositions and (ii) has the largest node ID, x_{\max} , among the nodes 14 in $\overline{V(P_0 \cup \dots \cup P_{i-1})}$. The largest node ID is used to select a single ear from multiple possible ears which each have the same nodal count.

At step 628, a left border node, e_i^l , is identified for the ear P_i 130 in the set $V(P_i) \cap V(P_0 \cup \dots \cup P_{i-1})$. At step 630, a right border node, e_i^r , is identified for the ear P_i 130 in the set $V(P_i) \cap V(P_0 \cup \dots \cup P_{i-1})$.

At decisional step 632, a determination is made regarding whether $d \in V(P_i)$. If $d \in V(P_i)$, the method follows the Yes branch from decisional step 632 to step 634. At step 634, the ear 130 starts on the d^l side, with $X(e_i^l) = X(d^l) = 0$. The ear 130 may also end with the destination node d 64 on the d^r side. This is the case when $e_i^l = e_i^r = d$. The nodes e_i^l and e_i^r are selected so that $X(e_i^l) < X(e_i^r)$.

Returning to decisional step 632, if $d \notin V(P_i)$, the method follows the No branch from decisional step 632 to decisional step 642. At step decisional step 642, a determination is made regarding whether node β is in $V(P_i)$.

5

10

15

25

port for the port 200 of node β leading to other nodes 14 in $V(P_i)$ is set to the red port. Returning to decisional step 653, if node β is not the right border node, the method follows the No branch from decisional step 653 and
5 returns to decisional step 618.

According to one embodiment, the coordinates for each $V(P_i)$ are made unique. For this embodiment, the coordinates of the j^{th} node from e_i^l , denoted as $n_{P_i,j}$ with $j > 0$, are assigned to:

10

$$X(n_{P_i,j}) = X(e_i^l) + \frac{M_i - X(e_i^l)}{C(P_i) - 1} \times j, \text{ with } j = 1, \dots, C(P_i) - 2 \quad (\text{eqn. 2})$$

$$M_i = \min_{n \in V(P_0 \cup \dots \cup P_{i-1})} \{X(n), X(e_i^l) > X(n)\} \quad (\text{eqn. 3})$$

15

where M_i is the smallest coordinate in $V(P_0 \cup \dots \cup P_{i-1})$ which is larger than $X(e_i^l)$.

20

FIGURE 7 is a flow diagram illustrating a method for reserving bandwidth for connectionless signals communicated between the nodes 14 in accordance with one embodiment of the present invention. The method begins at step 700 where a working bandwidth for a central node is reserved. A central node comprises any node 14 operable to receive traffic from a plurality of peripheral nodes.

25

At step 702, working bandwidth is determined for the central node based on the amount of working traffic that may be received from a particular peripheral node, assuming that the corresponding working path 104 is available. At step 704, protection bandwidth is determined for the central node based on the amount of protection traffic that

At step 706, the working bandwidth for the central node based on the working traffic from the peripheral node is compared to the protection bandwidth for the central node based on the protection traffic from the peripheral node. At decisional step 708, a determination is made regarding whether the protection bandwidth is greater than the working bandwidth. If the protection bandwidth is greater than the working bandwidth, the method follows the Yes branch from decisional step 708 to step 710. At step 710, additional working bandwidth is reserved for the central node in accordance with the difference between the protection bandwidth and the working bandwidth associated with the peripheral node.

Returning to decisional step 708, if the protection bandwidth is not greater than the working bandwidth, the method follows the No branch from decisional step 708 to step 712. At step 712, no additional working bandwidth is reserved for the central node based on the peripheral node.

From steps 710 and 712, the method continues to decisional step 714. At decisional step 714, a determination is made regarding whether there are more peripheral nodes that may contribute to the bandwidth requirement for the central node. If there are more peripheral nodes, the method follows the Yes branch from decisional step 714 and returns to step 702 in order to determine whether or not to reserve additional bandwidth for another peripheral node. However, if there are no more peripheral nodes, the method follows the No branch from decisional step 714 and comes to an end. In this way,

bandwidth is reserved for the central node in accordance with the bandwidth requirements contributed by each of the peripheral nodes from which the central node receives traffic.

5 According to one embodiment, bandwidth is reserved for the protection paths 100 and 102 to provide protection based on Quality of Service. The following notations are introduced to facilitate the discussion of this embodiment.

10 The set of all nodes 14 in the telecommunications network will be denoted as N_T . Each node 14 can be expressed as I_x where x is the index of the set N_T . The count of this set is $C(N_T)$.

15 The set of all unidirectional links in the telecommunications network will be denoted as L_T . Each member of L_T will be represented as l_i , with i being the index of the set L_T . The count of this set is $C(L_T)$.

20 Bandwidth reservation involves a determination of the required link bandwidth for protection purposes, expressed as P_{l_x} for any link l_x inside the network. In determining protection bandwidth reservation for a link l_x , traffic from other links and nodes 14 onto link l_x are considered, as discussed below.

25 For any working link l_y inside the network that sends working traffic onto link l_x , the bandwidth for this portion of the traffic on link l_x is expressed as $B_w(l_y/l_x)$ (bandwidth of the working traffic from l_y to l_x). If link l_y is broken, such traffic disappears.

For any broken link l_y inside the network, bandwidth for the protection traffic directed onto link l_x due to the broken status of link l_y is expressed as $B_p(l_y/l_x)$.

For any working node I_y inside the network that sends
5 working traffic onto link l_x , the bandwidth for this portion of the traffic on link l_x is expressed as $B_w(I_y/l_x)$ (bandwidth of the working traffic from node I_y to l_x). If node I_y is broken, such traffic disappears.

For any broken node I_y inside the network, bandwidth
10 for the protection traffic directed onto link l_x due to the broken status of node I_y is expressed as $B_p(I_y/l_x)$.

Thus, for protecting against a failure on link l_y , the protection bandwidth on link l_x , expressed as $P_{l_x}(l_y)$, is given by:

15

$$P_{l_x}(l_y) = \max(0, B_p(l_y/l_x) - B_w(l_y/l_x)) \quad (\text{eqn. 1})$$

Similarly, for protecting against a failure on node I_y , the protection bandwidth on link l_x , expressed as $P_{l_x}(I_y)$,
20 is given by:

$$P_{l_x}(I_y) = \max(0, B_p(I_y/l_x) - B_w(I_y/l_x)) \quad (\text{eqn. 2})$$

Thus, an array can be established by considering each
25 link and node 14 inside the network as follows:

$$[P_{lx}(l_1), \dots, P_{lx}(l_{C(L_T)}), P_{lx}(I_1), \dots, P_{lx}(I_{C(N_T)})] \quad (eqn. 3)$$

The array expressed in equation 3 can be numerically sorted and expressed as $\{p_1, p_2, \dots, p_n\}$ with $p_1 \geq p_2 \geq \dots \geq p_n$ and
5 $n = C(L_T) + C(N_T)$. For a single failure inside the network, the amount of bandwidth required for protection purposes is:

$$P_{lx} = p_1 \quad (eqn. 4)$$

10 For M failures, the amount of the bandwidth required for protection purposes is:

$$P_{lx} = \left(\sum_{j=1}^M p_j \right) \quad (eqn. 5)$$

15 Based on equation 3, it is also possible to specify a set of multiple failures for protection, depending on the protection policy (based on the importance of protecting certain links over others). Thus, this reservation mechanism minimizes the amount of bandwidth to be reserved
20 for protection purposes.

Although the present invention has been described with several embodiments, various changes and modifications may be suggested to one skilled in the art. It is intended that the present invention encompasses such changes and
25 modifications as fall within the scope of the appended claims.